ConneXions

The Interoperability Report

July 1990 Volume 4, No. 7

ConneXions— The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

ConneXions is published monthly by Interop, Inc., 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. 415-941-3399. Fax: 415-949-1779.

ISSN 0894–5926

From the Editor and Publisher

Yes, that's right, starting with this issue I have assumed both titles and responsibilities. If you're wondering if this will result in any significant changes in editorial policy, the answer is no. We will however slowly expand the scope of *ConneXions* to cater to our diverse audience, without of course losing the fundamental technical basis. Your comments and suggestions are welcome as always.

Our main feature this month is a report on the state of interior gateway protocols (IGPs) in the Internet. Currently there are two contenders for the position of standard protocol; OSPF developed by the Internet Engineering Task Force, and the OSI IS-IS protocol, extended to handle IP and well as CLNP packets. Daniel Dern spent a couple of months investigating the situation and interviewing key players in the community. The result, while not meant as a detailed treatment of the issue, should give you a good grasp of the technical and political climate surrounding this protocol development. Following the main article and the statements from users and vendors, is a brief overview of the working groups of the IETF tasked to focus their attention on various aspects of routing.

John Romkey is back again, this time with an article on the Packet Driver, a useful tool for DOS systems that allows several protocol stacks to access the same network interface card at the same time. Incidentally, John is working hard to develop the *Internet Toaster* which is due to be demonstrated for the first time during INTEROP® 90. More details about this project will follow in a future issue of *ConneXions*.

Also in this issue you will find a call for participation for the *SIGCOMM–SIGGRAPH Workshop on Graphics and Networking*, a book review and a letter to the Editor.

Next month we present another *Special Issue*, this time on Network Management and Security from a practical perspective. Stay tuned for this 40-page special report.

The INTEROP 90 Advance Programs have been printed and should be arriving in your mailbox any day now. Call us at 1-800-INTEROP or 415-941-3399 if you did not receive yours or need additional copies.

As is customary at the INTEROP conference, there will be a number of informal *Birds Of a Feather* (BOF) sessions, on topics related to networking. A dozen or so have already been scheduled and are listed in the Advance Program, but we invite anyone to suggest further topics for discussion. Call the numbers above and ask for myself or Susie Karlson to schedule a BOF. *—Ole Jacobsen*

Standards for Interior Gateway Routing Protocols

They're emerging, but still under debate and development

by Daniel P. Dern

Introduction

IGPs—*Interior Gateway Protocols*—have become a hot topic. Over the past year, IGPs have been the focus of a fair amount of debate and development, with the issues ranging from technical to operational to political. What's all the fuss about? Who cares, and why? Answer: Organizations building or expanding internetworks of host computers and/or LANs with multiple vendors, protocols, etc. care. Folks with the larger networks, like NASA, DoE, to be sure. Planners responsible for DECnet and IP architectures. Government agencies concerned about the soon-in-effect GOSIP FIPS.

Also, developers at router and host vendors. Planner, managers and users of regional nets and others, suffering from the limitations of RIP, the old IGP *Routing Information Protocol* (derived from the Xerox Network System's (XNS) routing protocol, also called RIP), bundled into many UNIX systems. And the *Internet Engineering Task Force* (IETF), tasked as they are with charting a course for Internet health and growth.

The IGP "question" is far too complex to completely address in one less-than-booklength article. What I intend to do here is:

- Briefly summarize key concepts for those readers not already up to speed.
- Attempt to state the "IGP question"—a meta-question, it turns out, as you'll see.
- Let spokespersons from leading networks, organizations, and vendors speak for themselves, presenting their current position and opinions.

And then Phill Gross, IETF Chair, will conclude with a recent policy statement.

What's interesting about the IGP discussions is that:

- Starting from the same base of information, different parties are picking different solutions, based on their point of view in the matter.
- Different people have different views on what the problem is, or which part of the problem they want to solve.
- We don't have clear consensus on whether the problem is predominantly technical, political or both.

Internet architecture

As we connect computers into networks, the information needs to know where to go, and how to get there. Inside a network, host computers may route information themselves; switches (e.g., packet switches) may otherwise do this. Networks combine into *internetworks*. *Gateways,* commercially called *routers,* are the junctions linking networks (WANs, LANs, sites, hosts, etc.), and forward packetized traffic from one net to another, using routing tables.

A collection of routers under a common administrative authority constitutes an *Administrative Domain,* or AD, previously called an Autonomous System, or AS). (Note: Multiple ADs could share the same underlying networks to interconnect their routers.)

Jack Haverty, Chief Architect, BBN Communications, recalls, "The notion of Autonomous Systems evolved in the 1980s for two reasons. One was to create an environment that could accept routers from multiple vendors with something more than static routing tables, and without locking everyone into a very rigid standard. The intent of this was to encourage the exploration of many different ideas by many separate groups, to get maximum research results.

Firewalls

The other purpose was to help us tackle the scaling problems of large and growing networks—how to carve them up into smaller sub-nets so that each was manageable. We wanted to define the place—the boundary—at which any particular Autonomous System's management could create a "firewall" to protect the activity within their system from disruptions caused by external problems. At the time, we didn't know what was the best way to create such a firewall, so the Autonomous System concept created the environment in which people could try a variety of approaches. (It's not clear that we yet know how to build strong firewalls.)"

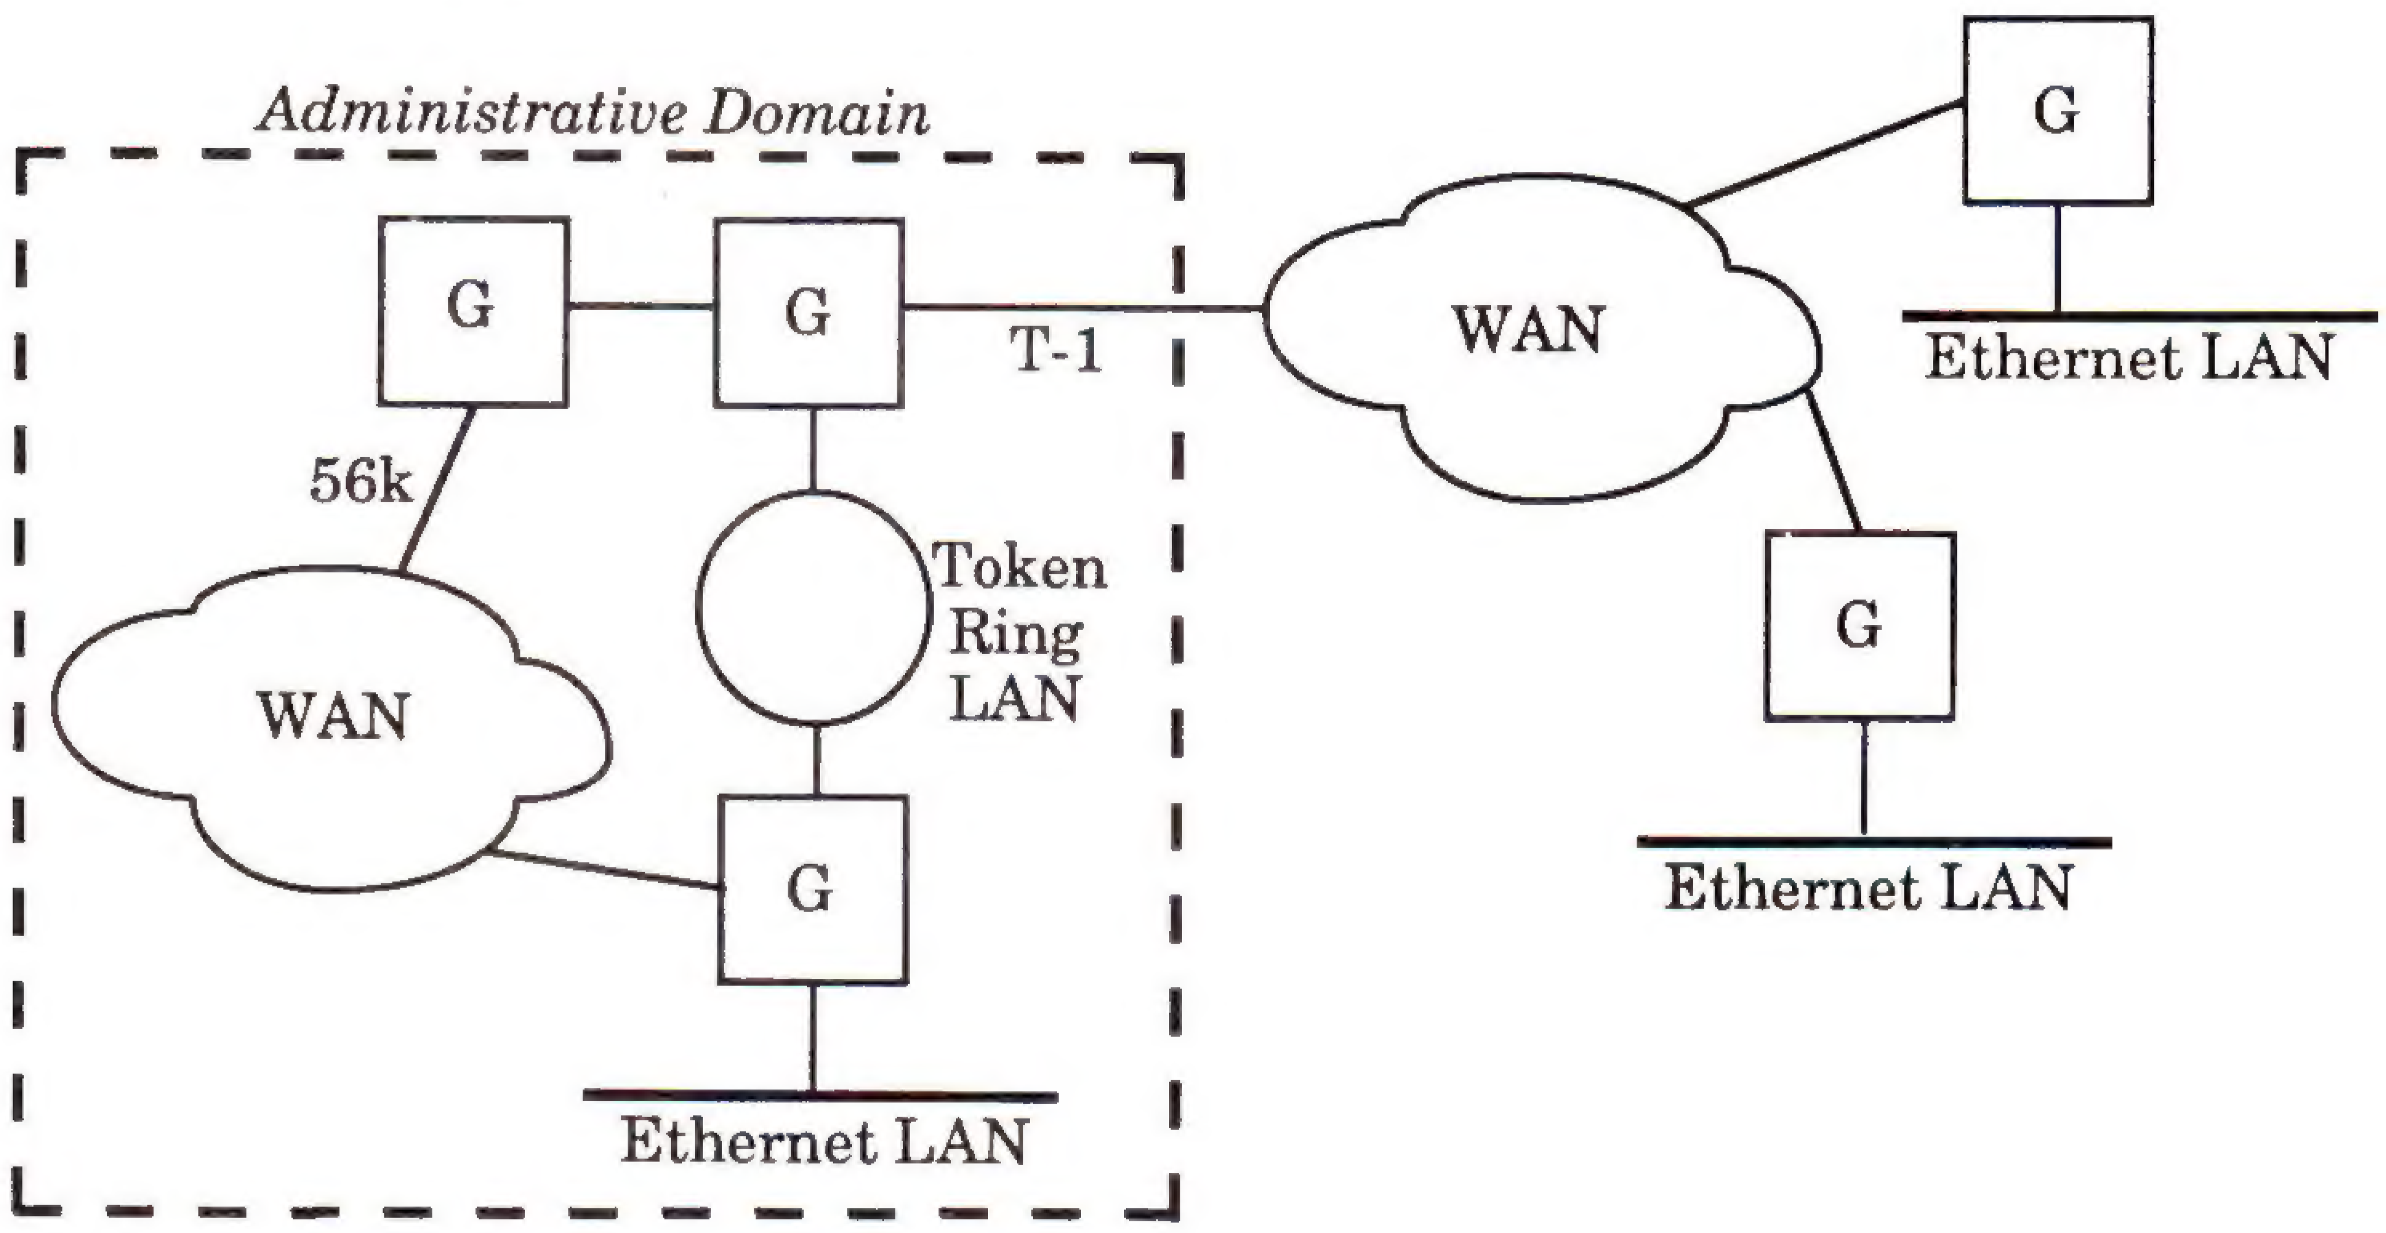

Figure 1: Internet architecture

Routing protocols

In addition to forwarding user traffic, gateways often exchange connectivity information, regarding what they see of network traffic and connectivity conditions.

This information comes from, and is sent to, other gateways using a commonly-spoken *routing protocol.* Gateway routing protocols are for gateway-to-gateway communications; they constitute separate, additional traffic from the data packets being forwarded.

Gateways within an Administrative Domain communicate via an *Interior Gateway Protocol,* IGP. (Gateways communicate between ADs via exterior gateway protocols, EGPs, now called *Border Gateway Protocols,* BGPs. (BGPs are outside the boundaries of this discussion.)

Dynamic and static routing

Each gateway stores a *routing table* which it consults to determine where to forward each packet it receives. In *static routing,* these tables are manually loaded/updated by network operators; whereas in *dynamic routing,* the tables are created and modified by the routing protocol algorithm running in each gateway on an autonomous, perhaps periodic basis.

continued on next page

Interior Gateway Routing Protocols *(continued)*

Dynamic routing offers significant advantages over static. Bob Hinden, Director of Internet Switching and Security at BBN, reminds us, "When there's network failure or congestion, dynamic routing picks the next route automatically. It lets you build networks that have alternate routes 'built in' to the topology, gives you self-configuring topologies that absorb changes and additions, and is intrinsic to reliability and availability."

The adaptive routing protocol used by the BBN packet switches in the ARPANET/Internet is a dynamic protocol. (By definition, there are no examples of "static routing protocols," as they have no algorithms, and no routing information traffic.)

Algorithms

Dynamic routing tables are calculated based on an *algorithm* built into the gateway code. The two dominant algorithm types are *distance-vector* (based either on "hops" or absolute distances), usually employing the Bellman-Ford algorithm; or *link-state*. Link-state approaches include *Shortest-Path First,* SPF (minimum time end-to-end transit); there are several SPF algorithms to choose from, usually employing the Dijkstra algorithm for calculating shortest paths.

Is one better than the other? Industry consensus appears to favor link-state. Many feel distance-vector algorithms won't handle today's bigger networks. "RIP's maximum hop count is 16, and that's too small for a large network," notes Vint Cerf, Vice President at the Corporation for National Research Initiatives. "And that's too small for a large internet."

Distance-vector algorithms such as RIP are also particularly vulnerable to 'black holes' in routing tables, where active loops cause connected portions of the internet to "disappear."

On the other hand, there are a lot of big distance-vector networks happily chugging away. cisco, for example, cites networks having hundreds of its routers, and 8 to 9,000 hosts.

Link-state does seem the contender, and SPF the preferred link-state method. "SPF isn't without its problems," notes Fred Baker, Director of Software Engineering, Vitalink. "But it tends to be a lot more stable, and tends to converge much faster. SPF is 'provably correct' in many cases where distance-vectors aren't provable."

Gateway information exchange

Every gateway needs an IGP. Many host and gateway vendors have implemented their own *proprietary* protocol, e.g., cisco, Vitalink, and Digital. There are also some *vendor-independent protocols*—RIP, OSPF announced by the IETF in October 1989 at the INTEROP conference, and the OSI IS-IS (*Intermediate System-to-Intermediate System*) IGP.

In an internet, it's possible—in fact, more than likely—that there will be gateway/router boxes from more than one vendor. Haverty notes that users often find they already have a multivendor inventory. "The question becomes, how best to utilize it."

One common protocol?

Gateway routing information needs ultimately to be exchanged via a common-spoken protocol; likewise, routing calculations need to be done for each protocol.

These are related, but separate, questions. Which protocol? That depends. Many router vendors have their own proprietary protocol. An all-cisco internet can speak IGRP—cisco's proprietary protocol—, for example. But what of the internetwork with multiple vendors' routers?

Commonly-spoken protocols are needed. Many vendors, such as cisco, Proteon and Wellfleet, support multiple protocols. RIP was created as a vendor-independent standard, bundled into UNIX systems. More recently, OSPF was created to provide a high-functionality vendor-independent IGP, so that multivendor networks would not be forced into lowest-common-denominator service.

The next killer question: suppose your network needs to support more than one routing protocol at a time? More precisely: suppose your network has more than one type of traffic, e.g., TCP/IP, DECnet Phase IV or V, Novell, XNS, AppleTalk, and OSI packets?

There are several ways to solve this—each with its own set of tradeoffs in development effort, device cost, control traffic, and operational impact. And each approach has its advocates and critics.

Parallel stacks

"Parallel stacks" running in the routers is one approach. Here, each protocol runs its own algorithm, in turn consuming some degree of CPU cycles for calculations, and memory for code and table storage. For instance:

- Process A – IP routing
- Process B – DECnet/OSI Phase V routing
- Process C – yet another routing algorithm

Baker suggests that the code and data for each additional protocol added to a router will need in the neighborhood of 30–50Kbytes of additional memory, which he translates into around $500 in manufacturing costs—probably adding $1,000 to $2,000 to the final price of each router.

The advantage: relatively straightforward, modular software engineering. Each protocol has its own resources; they co-habit easily in the box.

Disadvantages: One, device cost. The minimum just got raised. The low-end cost of routers determines who can get into the game, as niche vendors, and as users. Two, higher traffic loading...maybe.

Overhead versus performance

Every algorithm needs its periodic fix of "network news." The more protocols running, the higher the volume of non-user traffic—and line costs tend to be the dominant cost of network operations.

In a stable multiprotocol network, meta-traffic shouldn't impose a significant burden—but the network-wide reaction to a major event could conceivably saturate links with control traffic, leaving no bandwidth for user traffic. There are steps that can be taken to minimize this effect however.

How often will net-throttling events happen? "That's a hard question to answer," William Seifert, Wellfleet's VP Strategic Planning and Chief Technical Officer. "It's a direct function of scale. If your network is, say, the size of 3M's, it's likely that something will be changing in any given instant. In any hour, you're bound to have change."

continued on next page

Interior Gateway Routing Protocols *(continued)*

This is a mostly theoretical problem, acknowledges Seifert—at this point. "Nobody has observed these problems yet. Because there isn't enough deployed to have seen it. But problems of internetworking have always been addressed in this way—because we can't afford to ignore them. The Internet was created under this philosophy. And we know enough that we assume such problems will happen, and we want to prevent them."

Ships In the Night vs. Integrated Routing

"The way that most vendors are currently approaching the problem is what's come to be called *Ships In the Night*," says Gross. Info packets from various IGPs elbow past each other like ships passing in the night, giving rise to this approach's nickname.

"In S.I.N. implementations, everything runs as totally independent, parallel vertical stacks. So if you have a router doing DECnet, IP, OSI and XNS, you would be running in this case four different routing protocols, each with its own algorithm.

Suppose the routing data was somehow integrated into a single set of messages. Could a single routing protocol calculate routes for use by multiple protocol stacks—say, for IP and OSI? Pros: less traffic—probably. Cons: possible compromises in usefulness of information content. Specific concern: OSPF control traffic cannot be modified to also serve OSI IS-IS without severe surgery, due to essential differences such as address length. The OSI IS-IS protocol, on the other hand, can be extended to provide information to IP."

"In Dual IS-IS routing," says Lyman Chapin from Data General, "the ANSI/ISO standard IS-IS routing protocol is used to distribute intra-domain routing information among intermediate systems (gateways), and being able to use that information to do routing, for either DoD IP, or ISO IP."

In still further integration, one algorithm might—possibly—be extended to generate routing tables for multiple protocols. "This approach also means less storage requirements, for only one set of routes," notes Radia Perlman, Digital Equipment Corporation. "And less control traffic, because you only have one routing algorithm running."

Would multiple protocols be 'resource pigs' in operational nets? Suggests Gross, "If you have a 100-node network, all running all protocols, then an integrated approach may be more attractive. But that's not the case at present. Most networks are imbalanced—say, mostly IP with a little DECnet or OSI. Then you don't get the same overhead. You're sending a lot of IP to a lot of nodes, a little OSI to a few OSI nodes, etc. And a few nodes doing both."

An IGP for the Internet

The IETF would like to standardize an IGP routing algorithm for IP. Lack of one is a source of continual concern to those in charge of networks today, and those planning what lies ahead. (See separate comments from Philip Almquist, Stanford/BARRNet, and Milo Medin, NASA/Ames.) Urgent needs perceived by the community include:

- A robust, flexible protocol that will run in their multivendor environments.
- Cohabitation for IP and DECnet traffic.

"It's sort of amazing that the Internet has survived as long as it has without an official one, if you don't count RIP," comments Perlman. She was the architect for DECnet's routing layer for a long time, and has helped design the Dual-routing IGP.

The vendor community wants to address this need. At the same time, outside the Internet, and working its way within, is a growing interest in OSI standards.

Within the link-state world, there are two contenders for IGP: The Open SPF Protocol for IP (OSPF), and the IS-IS ISO protocols for OSI traffic. There are also two camps on approach: *multiple stacks,* or *integrated routing,* also known as Dual-stack.

Dual Stack IS–IS routing

The IETF has a working group, chaired by Ross Callon of DEC, vice-chaired by Steve Willis of Wellfleet, that has been defining a way to convey IP IGP information via extensions to the ISO IS-IS protocol.

"The IS-IS working group is convinced that the IS-IS protocol can be augmented to do that job, to distribute routing info for DoD routing," says Chapin. Callon adds, "The previous incarnation, which is a subset of the IS-IS protocol, has been running on NSFnet for over two years."

No easy choices

Gross points out, "This is a very complex and difficult technical issue. There are substantial differences among the various sides and approaches—and they all have their merits."

Steve Crocker, Trusted Information Systems, offers his thoughts: "Routing is an area of some technical risk, in that if you do it wrong it has dramatic impact on the network, in terms of serious disruptions. So one thing I'd like to see is a thorough understanding of the technical issues, making sure these protocols are robust."

The basic difference of opinion is some people believe there is substantial benefit to the Dual IS-IS, because it gets people working on OSI faster than otherwise, and these people believe that working on OSI is good in itself. Others believe that nothing which is sub-optimal is acceptable to the Internet. (See Medin's comments.)

"There are two orthogonal issues to remember," says Perlman. "One, technical comparisons of OSPF versus IS-IS—which I feel are nearly identical—and two, is it better to have only one routing algorithm, which saves the design and development effort to implement a second routing algorithm, and minimizes CPU and bandwidth consumption. And since it's easier to modify IS-IS to support TCP/IP, IS-IS makes the most sense."

Parallel stacks may make for a very busy router, comments Gross, "But it is a 'no brainer' approach. You don't have to do anything complicated. It's in some ways the trivial and obvious method. You're running multiple virtual routers, all in the same hardware. And it should be pointed out that this approach is available from most gateway vendors now."

Many network managers feel this needs less maintenance, he notes. "Some protocol designers believe that integrated routing is a very clever approach. But a large group of network operators and managers prefer the reverse. They like segregation of functionality, it's easier to track down errors and problems."

continued on next page

Interior Gateway Routing Protocols *(continued)*

However, Gross adds, Dual IS-IS means that the people managing the network only need to manage one coherent routing protocol. "It means you'll only have one that will require upgrades in the future."

For many, however, the future is a long way off. "OSPF has a public domain implementation which is available today," noted Gross in March 1990. "The Dual IS-IS probably won't exist for at least another six months. That doesn't seem like a long time, but there is a lot of pressure in the Internet world for something right now."

OSPF availability

Hinden adds, "I think the current Internet consensus is that OSPF is farther along than a Dual IS-IS—it's been implemented by Proteon, it's available as a public domain UNIX implementation, it has been deployed in some regional nets, and in the NASA backbone, which has a lot of Proteon routers. So we'll soon have serious operational experience with OSPF. The pure ISO version of IS-IS is now very well defined. It's a draft status in ISO—but the IP extensions are new. So it's about a year behind OSPF."

In a message to the Internet TCP-IP mailing list, Haverty gave his thoughts on the IGP question. (See separate text following main body of article.) "The gist of my message," Haverty noted in a follow-up message (excerpted and paraphrased with permission), "was: Everybody asks whether OSPF or IS-IS is the better solution. I asked the question, could somebody please tell me what the routing problem is, at the level of not just 'connect everything together'?

Is it the problem to create a routing mechanism which will respond to transients within seconds? Or is it to create one that operates with minimal overhead? Or that will operate in the presence of failures? Or which is very simple to implement and test?

These are all different problem statements, which in some cases are contradictory in terms of solutions. I was asking somebody to please 'state the problem,' because I'm not sure everyone is working on the same problem. Both (all) sides may be right—because they are working with different goals."

Haverty adds, "My experience has been that whenever you see strong argument and disagreement on technical issues, it often means the different camps are really working on different problems.

Motivations

The regional nets have a multivendor router environment, so they only can run RIP, which often provides inadequate service. This accounts for much of the desire for a near-term solution: to replace RIP.

By no coincidence, algorithms in the OSI's IS-IS and the one in DECnet Phase V are nearly identical—DEC contributed their work to OSI, and have kept up with any modifications.

In terms of the Internet community, OSI is coming, and TCP/IP will be around for a while—wouldn't one routing algorithm be better than two—and IS-IS can be made to do the trick.

IESG position

At the February 1990 meeting of the *Internet Engineering Steering Group* (IESG), Chair Phill Gross made the following recommendation to the *Internet Activities Board* (IAB):

> "There is a pressing need for a high functionality *open* Intra-AS [AD] Interior Gateway Protocol (IGP) for the TCP/IP protocol family. Users and network operators have also expressed a strong need for routers from different vendors to interoperate....
>
> The IESG, reflecting the discussion in the IETF plenary at Florida State University, decided that both protocols [OSPF, and ISO IS-IS enhanced to support IP in tandem with CLNP] need substantial operational experience before either could be made full Internet standards or recommended to the IAB as the 'Recommended' IGP for the TCP/IP protocol family.
>
> The practice within the IETF has been to allow a protocol to begin the standards process with a designation of Proposed Standard prior to gaining field experience. Extensive field experience is required prior to an advance to Draft or Full Internet Standard.
>
> Therefore, the IESG recommends that OSPF be designated a Proposed Standard at this time. Further review and advancement as an Internet standard will await the outcome of current ongoing field trials.
>
> The IETF and IESG have expressed interest in the integrated routing that is promised by the Dual IS-IS, but also expressed concern about potential complexity and side-effects."

Conversion rarely happens overnight

Even if adding OSI (and subtracting TCP/IP?) from Internet operations starts soon, it won't be over for five to ten years. Rolling one protocol out and another in rarely happens in a simple, late-night "throw the switch" changeover. The NCP-to-TCP changeover in the ARPANET during the early-80s, for example, took a lot of coordinated scheduling, late nights and follow-through, according to Interop's Dan Lynch, who assumed responsibility for much of the site and host coordination. "And this changeover only involved about 200 hosts," adds Cerf. Migration will need careful planning, and won't happen overnight. Many people want to preserve their current investment and technology, and the migratory path is non-trivial.

Remarks Gross, "One lesson we learned very painfully in the CMOT versus SNMP controversy is that operational experience and field testing are *very* important. The Internet approach is 'prototype, experiment and field test.' We want to do this for routing. I'm enthused about integrated routing, as a notion. But I'll be more enthused when there's an implementation, and experience...with big networks that have both types of nodes.

Chapin at Data General is in favor of OSI and Dual IS-IS for two pragmatic reasons: "I find that whenever you try to do two difficult things instead of one, you end up doing them less well. Also, there's a finite pool of people with experience in connectionless networks. Getting this IP community involved in the problem of effectively and efficiently deploying OSI protocols would shorten the cycle enormously."

The next release of Berkeley UNIX will include OSI, thanks to contributions such as transport and network layers from the University of Wisconsin, and upper layers from ISODE.

continued on next page

Interior Gateway Routing Protocols *(continued)*

Will this encourage OSI applications and OSI transport—leading to faster deployment of OSI traffic on the Internet?

As should be evident, there are a lot of top-level people thinking deep thoughts on the matter. The author finds it unlikely that the IP network planners will accept sub-optimal solutions casually—and if that to some extent determines their vendor selections, so be it. Equally, vendors will have to decide their strategies, based on sales prospects.

I find the arguments concerning higher hardware costs, and even those regarding traffic requirements, less than compelling—because hardware costs have been dropping, and bandwidth growing. Will control traffic loading be significant in a T-1 pipe?

And so far, it seems that the multi-integrated question only applies to TCP and OSI? If integrated routing doesn't deal with any other protocols, then we've only reduced an *n*-protocol problem by 1. Also, remember that OSI *per se* is not sufficient for all DECnet requirements.

What lies ahead

Network needs and available commercial offerings are ripe for the first rounds of major deployment of multivendor IP networks, and forays into multiprotocol networking. At the same time, there are a number of major yet-to-be-solved research questions into the fundamentals of large networks.

"We'll see more research into complex environments with multiple policies, administrations, management structures—and how to make these appear as one seamless service to the end users," suggests Haverty. "It's still not clear how to do this, or even what all the questions are. We also will see the next-generation Exterior Gateway Protocol—we've known from the start our initial efforts would need replacing. One candidate is BGP—the Border Gateway Protocol.

Pragmatically, we should continue deploying operational systems soon, based on current implementations, so we can accumulate real-world experience to complement the research thoughts on how to solve unsolved problems."

Acknowledgements

Many people have been very patient and helpful, in what turned out to be a much larger project than anticipated. My thanks to all who helped, and apologies to any I've inadvertently misrepresented. In addition to all people mentioned or quoted in the article and related text, thanks are due to Jim Herman, Vint Cerf, Craig Strauss, and Sheryl Schultz. And last but hardly least, to Phill Gross, who (innocently) started this with his phone call late one Friday afternoon.

Ed.: Statements from users and vendors follow on the next page——>

DANIEL P. DERN is a Watertown, Mass-based free-lance writer specializing in technology, science and industry. A frequent contributor to *ConneXions*, including last year's ARPANET historical retrospective, Dern writes for leading publications and vendors in the network and computer industry, as well as writing humor columns, science fiction, and musical theater. He was previously PR Manager at BBN Communications. He can be reached at `ddern@world.std.com` (Internet) or `dandern` (MCIMail).

Statements from users and vendors

[Author's note: Quotes and information are based on discussions held up through April 1990, and may not reflect newer announcements.]

User: Milo Medin
Network Architect
NASA Science Internet Project Office
NASA Ames Research Center

Medin is responsible for the networking needs of NASA scientists in the US and worldwide, notably Pacific Rim backbone connections to the Internet. Their current network has forty (Proteon) routers, supporting DECnet and IP traffic, and connects both sites and regional NSF subnets. Medin estimates packet traffic is more or less evenly mixed among interactive remote login, file transfer, and e-mail, including "a lot of supercomputer access, and big files like LANDSAT images and the Voyager Neptune encounter photos." They also support programs like the Hubble Space Telescope and Galileo. Milo has been a major contributor to the OSPF project, drawing on his experience in the requirements for large operational networks.

"I strongly disagree with those who claim OSPF and Dual IS-IS are too close as to make it a political debate. OSPF is optimized for IP; IS-IS is built for OSI. My concerns include:

- Issues with the base IS-IS protocol.
- Issues with the Dual IS-IS approach as currently defined by the IETF working group.
- Issues with the general approach of one routing protocol for multiple protocol suites, versus Ships In the Night.

Choice of architecture

The Integrated camp has two strong statements: That you should use an IS-IS protocol for IP, and that you should run an *integrated* IS-IS, for OSI CLNP and for IP. This means you are not just making a decision on routing protocol, but also on your whole architecture. They are suggesting we pick an architecture which we have little experience with, rather than S.I.N. which vendors and users have ample experience with.

You can argue why IS-IS is a good way to route IP. But adding the integrated routing baggage is, I believe, the wrong approach. I don't believe it will make the coding easier for router vendors. You'll have bigger code with more interdependencies—which is totally against the modular philosophy. Two separate stacks will be easier to support.

There are differences in the protocol structures, of course. My point is, the Dual implementors don't want to modify IS-IS to optimize IP routing flexibility and performance.

Concerns

My concerns (needs which OSPF solves) with the Dual IS-IS, integrated routing, and the current proposal as it stands include:

- Level 1 information
- Virtual links for AS robustness
- Variable length field inefficiencies
- Authentication
- Metric restrictions and suboptimal compromises

continued on next page

Statements from users and vendors *(continued)*

- Area partitions and healing
- Black holes in areas
- IP Type of Service issues
- Management, implementation, operations and support concerns
- Problems introduced for OSI deployment

Negative benefits

And further, I don't believe things like metrics and areas should be the same for OSI and IP. That's the central issue for me: Are you trying to do the absolute best possible job of supporting IP? Or to gain the benefit of integrated routing? And if so, does this override what I lose, in terms of existing IP optimizations? Why should I give up OSPF for something that gives me 'negative benefits'?

Also, remember that nobody in the Internet today is passing operational OSI traffic. It's going to be a year or so before these new protocols are working. OSPF is available now, and many of us in the Internet need it *immediately*. I can't wait.

Our operational networks have *already* outstripped the capabilities of existing routing protocols. That's why the IETF started the OSPF group 2 years ago, because it was becoming clear we would have to get better protocols just to survive.

OSPF already running

In addition, the NASA Science Internet converted to OSPF on Friday, April 13, 1990, and it's working very well. The amount of routing traffic has dropped drastically from what RIP required, and yet the network reroutes around down lines *much* more quickly, using much less traffic. We switched off a main trunk in the network, and 3 seconds later, the entire system had rerouted. And the external route tagging has greatly simplified our network configuration tasks. In short, we are using it in a demanding *production* environment, and have been for a few weeks, and it's working extremely well."

User: Philip Almquist
BARRNet

Almquist is chair of the technical committee for BARRNet, active in the Stanford network, and designed the networks for the 1988 and 1989 INTEROP shows.

Severely hampered

"BARRNet is severely hampered by the present state of IGPs. The INTEROP networks were equally constrained. For the INTEROP networks, the problem was that we were using routers from multiple vendors. The only IGP they all had in common was RIP.

RIP in general has no authentication; because UNIX systems commonly can 'speak RIP'—and as as a rule, by default, boot-up includes a RIP daemon. And this daemon may be using totally meaningless information, like an ARPANET host address even though it isn't on the ARPANET at the time.

RIP unusable for INTEROP

The backbone for the INTEROP 89 network was multivendor, including equipment from cisco, IBM, Proteon and Wellfleet. The potential for bogus behavior made RIP unusable, because of the impact it would have on problem (creation and) resolution. So lots of the network was forced to use static routing—lots of manual entry, often by sleep-deprived staff, and zero robustness.

We don't use RIP in the Stanford network. We don't need to, because at present we have a single-vendor router environment. BARRNet has other concerns, but equally real. We have routers from multiple vendors. RIP does not have a good way to keep track of where information comes from, e.g., NSFnet routes, BARRNet-internal routes, MILNET routes. This severely limits flexibility for alternate routing."

Q: What do you want? Specific solutions, or something soon?

Standard needed

I think we're hurting more from doing without than we would from what might turn out to be a less-than-optimal solution. Part of the reason we want a standard is that we cannot replace large amounts of installed hardware. If you can enforce vendor selection, you can get acceptable routing today without there being a standard. But for examples like BARRNet, or the INTEROP network, you have to compromise and sacrifice a lot you really would prefer not to."

Host vendor: Data General Corporation
Lyman Chapin
Senior Consulting Engineer
Network Systems Development Division
Software Development

"We have a line of UNIX workstations and servers, which offer TCP/IP services, so we need IP. We also need OSI, or at least a statement of what we'll be doing, for all government bids.

OSI routing protocols only

We're only interested in OSI routing protocols. If we were a router vendor only, we'd have more motivation to support other protocols. But we're not a router vendor; we supply software for our AViiON servers to enable them to act as either DoD or ISO IP gateways. But that's not our principal business. So we feel comfortable concentrating on IS-IS.

Internet engineers needed for OSI development

One reason I support the Dual IS-IS I feel there's a limited pool of experienced talent available available in the engineering community. OSI needs the people with experience in connectionless networks, and they're mostly working in the IP world for now. We need these people in OSI, or it will take another decade for the OSI world to develop its own expertise. Bringing the IP community over will shorten this cycle tremendously."

Host vendor: Digital Equipment Corporation
Radia Perlman and Ross Callon

DECnet/OSI Phase V uses the SPF-based OSI IS-IS protocol.

Perlman: "I suspect Digital prefers the Dual-routing approach—implementing one routing algorithm which supports both TCP/IP and OSI, as least costly, but they will do what they need to. If the TCP/IP community says you have to use TCP/IP-derived algorithms, we will."

Coexistence and interoperability

Callon: "We are putting very high importance on the coexistence and interoperability of three things: Phase IV, Phase V—which is OSI-based—and TCP/IP. Digital will support IP through our routers (gateways), which can communicate directly IP systems, to VMS systems, etc.

continued on next page

Statements from users and vendors *(continued)*

Digital multiprotocol routers will need to be able to forward IP packets to destinations reachable via other external IP routers (using EGP, RIP, OSPF, or whatever). There is a clear need for application-layer gateways, to allow IP end systems to talk to DECnet/OSI end systems. For example, e-mail gateways. Both Ultrix and VMS systems will need multiprotocol support, so that each end system can originate and receive both IP and DECnet/OSI traffic.

Configuration easier with one protocol

Probably the most important advantage of the integrated IS-IS approach is the ease of network management and configuration. With the integrated approach, you have only one coordinated routing protocol to configure. With the Ships In the Night approach, you have two or more routing protocols running simultaneously. Thus you need to separately configure each, and problems with one protocol may create problems with the other(s), potentially resulting in problems which are difficult to debug."

Host vendor: IBM Corporation
Yakov Rekhter
Research Staff Member
T.J. Watson Research Center

Rekhter was involved as a representative from IBM in the design architecture and implementation of routing for the NSFnet backbone. The following opinion does not necessarily reflect any position of IBM.

Dual Stack better

"IBM has very strong interest in promoting standards, like TCP and OSI. For IS-IS, my personal opinion is that the Dual protocol is the best technical solution. There may be minor issues why S.I.N. (Ships In the Night) is better, but from a technical point of view, I believe Dual-stack is the best solution.

OSI is coming, like it or not. So we should make the best of it. And make the coexistence with TCP/IP as easy as possible. With OSI, both hosts and gateways will become more complicated, supporting two protocols. You don't want to add additional complexity. Dual IS-IS is easier than S.I.N. to manage—it's a single protocol.

IS–IS lowest denominator

From a vendor point of view, I can implement whatever customers will pay for. But IS-IS is the lowest denominator. It has to be done, for OSI. If you do it right, you also can support TCP/IP, whereas OSPF can't handle OSI."

Router vendor: cisco Systems
Douglas Tsui
Manager, Product Marketing

cisco Systems' routers support RIP, DECnet, OSI, XNS, IPX, AppleTalk, X.25 and others. IGRP, cisco's proprietary IGP, is distance-vector based. They have announced support for DECnet Phase V, will support OSI IS-IS, and "may add OSPF."

OSPF not proven yet

"We feel our algorithm has proven itself very well in large networks. It recovers very well. We feel there is room for both approaches.We think that OSPF is not at this time proven in very large networks. We will see what happens.

We expect to support RIP, IGRP, maybe OSPF, and IS-IS for OSI. We currently can handle IP over these routing protocols, and also do OSI over IGRP. Most other OSI implementations have to use static routing today. We can do dynamic."

Router vendor: Proteon Inc.
John Moy
Senior Staff Engineer

Proteon's router family, the P4100 Router and P4200 Router, support protocol stacks including IP, DECnet (IV), AppleTalk, XNS, IPX (Novell NetWare), Apollo Domain, and OSI. TCP/IP routing protocols supported include OSPF, RIP, EGP. Proteon does not employ any proprietary protocols, prefers link-state, and currently implements completely parallel stacks. Moy was a member of the IETF committee which defined OSPF.

"We've been a multiprotocol routing company for four years, which is as long as anyone's been in the business. We're strong proponents of implementing a completely parallel stack, which is what all multiprotocol routers currently have implemented.

We believe that if you have a separate network protocol, it should run a separate routing protocol with its own update traffic. This allows you to tailor the routing protocol to the particular needs of the protocol stack. It also makes network problems easier to debug, because problems in one protocol stack will not transfer to other stacks.

Parallel stacks more flexible

Parallel stacks give much more flexibility. With two separate routing protocols, you don't tie your OSI policy to your their TCP/IP policy. E.g., you may want to run IP in certain parts of your internet, and OSI elsewhere. Using separate routing protocols, your IP and OSI areas don't have to match or intersect, which they would if you ran a single integrated routing protocol. Yes, multiple stacks may mean more development effort, and take up marginally more system resources—but we think the benefits override these issues."

Router vendor: Vitalink
Fred Baker
Director, Software Engineering

Vitalink's TransPATH routers support its proprietary SPF-based TransPATH protocol, as well as other dominant stacks.

"Vitalink is in a somewhat unusual position, because we use a mix of both bridges and routers, in building networks.

As for 'which protocol,' we probably echo Dave Oran at Digital—that if all you are doing is IP, you need an IP-specific protocol. If you're building a multiprotocol router and a multiprotocol network—and the biggest reason Vitalink sells its stuff is for multiprotocol nets—then minimum event-based administrative traffic when an event happens is the fundamental reason for a using a single integrated algorithm.

continued on next page

Statements from users and vendors *(continued)*

OSPF and IS–IS similar

OSPF and IS-IS are extremely similar, in terms of how they do SPF. Doing their various route calculations would usually yield the same answers. So why not extend IS-IS instead of adding network load?"

Router vendor: Wellfleet Communications
William Seifert
Chief Technical Officer

Wellfleet router products include their FN Feeder Node, LN Link Node, and CN Concentrator Node. Protocols carried include TCP/IP, DECnet Phase IV and DECnet/OSI Phase V, XNS, Novell IPX, AppleTalk, and X.25. Wellfleet does not have a proprietary IGP. Wellfleet currently implements completely parallel stacks; TCP/IP IGPs include RIP and Extended RIP, with OSPF scheduled for late 1990.

"In the long term—which I think means sometime around 1992 and after, maybe sooner—we're headed is Dual-stack environment. IP routing, OSI routing, all in one piece of hardware, same sets of wires. For vendors, the question is how to support this and accommodate this within router products, and how to reduce this model's impact in terms of network resource utilization.

The market prefers OSPF

In principle, that means issues like internet responsiveness, bandwidth and circuit consumption, manageability, etc. Wellfleet feels that the market has spoken, and they have said: OSPF. It's *also* clear to us that IS-IS will be a necessity for OSI environments. We'll see TCP/IP around for a while. The GOSIP won't force existing users to change... But OSI is the direction...

Also, don't forget that a true DECnet Phase V router has to be more than just an OSI router. You also have to be able to route OSI traffic from non-DEC end systems, and also DECnet Phase IV end system traffic, such as PDP/11s, which can't migrate to Phase V."

Router vendor: 3COM
Clint Ramsay
Product Manager, Bridges & Brouters,

3COM routers include the BR/2000 Brouter and the BR/3000 Remote Brouter. 3COM is a standards-based company, routing XNS, TCP/IP, and OSI. Support for all the major stacks: XNS RIP, TCP/IP RIP, OSI ES-IS (static and prefix routing), and will support OSPF and OSI IS-IS in the next release. They are currently evaluating Dual IS-IS.

Parallel stacks most effective

"We've been selling routers since 1983—we've most recently introduced our family of multiprotocol brouters. Our position is that the world will continue to use multiple protocols, and that completely parallel stacks certainly at present is the most effective way to implement and support them. Particularly when you start to consider policy-based routing."

———*end*———

IETF Working Groups on Routing

The Internet Engineering Task Force (IETF) has several working groups that tackle routing protocol issues. These are listed below.

The IETF *IS-IS Working Group* will develop additions to the existing OSI IS-IS Routing Protocol to support IP environments and dual (OSI and IP) environments.

The *Interconnectivity Working Group* is working to develop the *Border Gateway Protocol* (BGP) and BGP technical usage within the Internet, continuing the current work of the Interconnectivity Working Group in this regard.

The *Multicast OSPF Working Group* will extend the OSPF routing protocol so that it will be able to efficiently route IP multicast packets. This will produce a new (multicast) version of the OSPF protocol, which will be as compatible as possible with the present version (packet formats and most of the algorithms will hopefully remain unaltered).

The *Open Systems Routing Working Group* is chartered to develop a policy-based AS-AS routing protocol that will accommodate large size and general topology.

The goals of the *Public Data Network Routing Working Group* are the development, definition and specification of required routing and gateway algorithms for an improved routing of Internet datagrams through X.25 Public Data Networks (PDN) to allow worldwide inter-operation between TCP/IP networks in various countries. In addition, the application and/or modification of the developed algorithms to interconnect local TCP/IP networks via ISDN (Integrated Services Digital Network) will be considered.

For more information about how to participate in these groups and in the IETF in general, send a message to:

`ietf-request@venera.isi.edu`

Special Routing Issue

To learn more about routing, read *ConneXions* Volume 3, No. 8, August 1989, *Special Issue on Internetwork Routing*. Still available for only $15, this 56-page report contains many articles on TCP/IP and OSI routing protocols written by experts in the field. Order yours today! All other back issues of *ConneXions* are also available for purchase. Call us at 415-941-3399 or 1-800-INTEROP and ask for our free index pages to pick the issues you want.

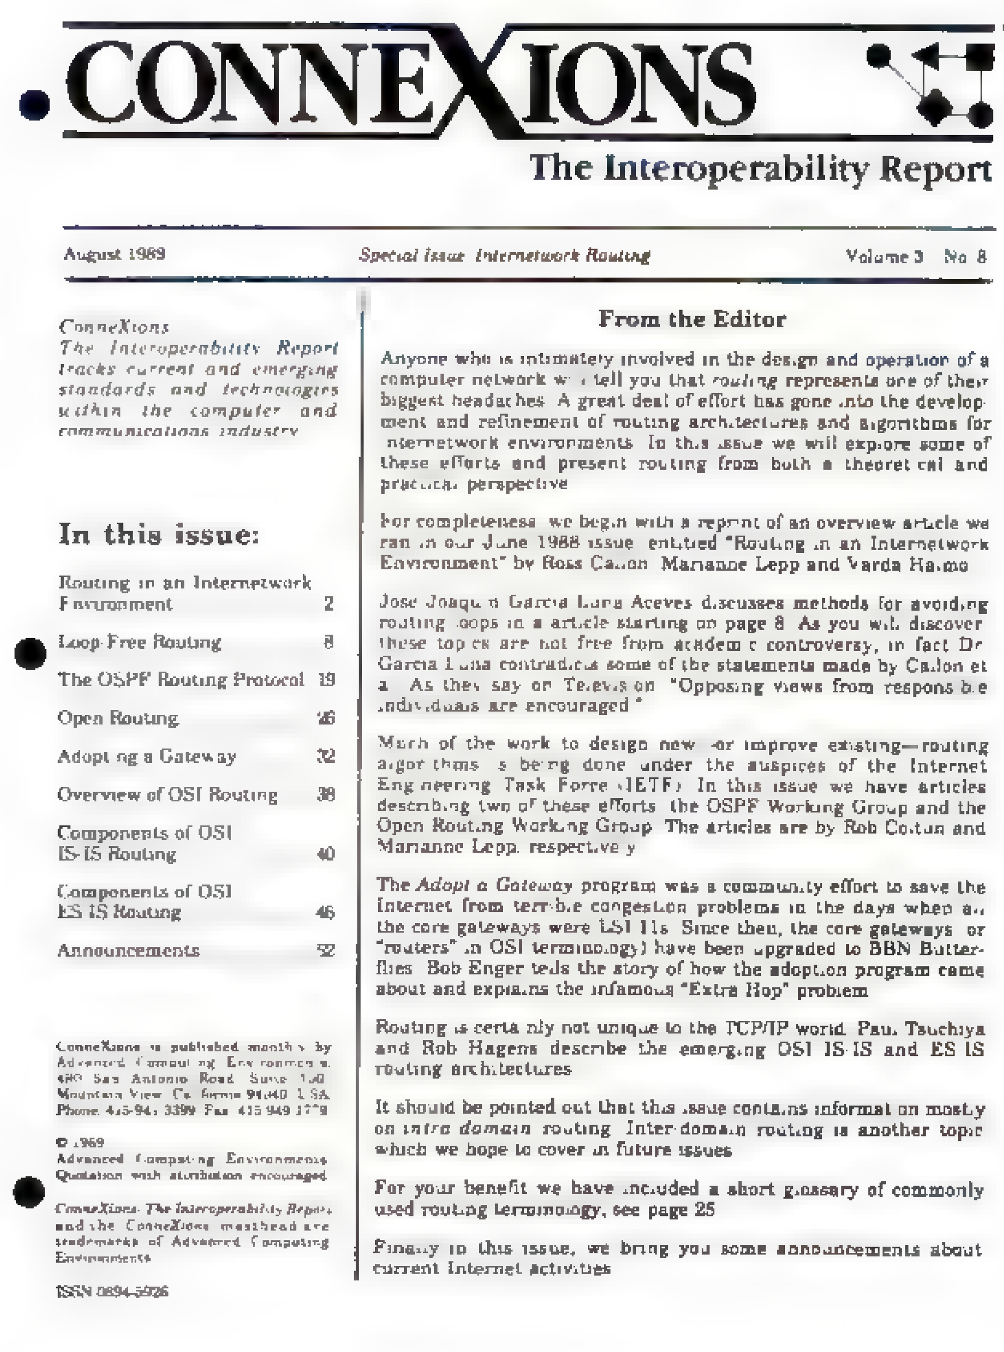

CONNEXIONS
The Interoperability Report

August 1989 | *Special Issue: Internetwork Routing* | Volume 3 No. 8

ConneXions The Interoperability Report tracks current and emerging standards and technologies within the computer and communications industry.

In this issue:

Routing in an Internetwork Environment	2
Loop-Free Routing	8
The OSPF Routing Protocol	19
Open Routing	26
Adopting a Gateway	32
Overview of OSI Routing	38
Components of OSI IS-IS Routing	40
Components of OSI ES-IS Routing	46
Announcements	52

ConneXions is published monthly by Advanced Computing Environments, 480 San Antonio Road, Suite 100, Mountain View, CA 94040, USA. Phone 415-941-3399 Fax 415-949-1779

© 1989 Advanced Computing Environments. Quotation with attribution encouraged.

ConneXions–The Interoperability Report and the ConneXions masthead are trademarks of Advanced Computing Environments.

ISSN 0894-5926

From the Editor

Anyone who is intimately involved in the design and operation of a computer network will tell you that *routing* represents one of their biggest headaches. A great deal of effort has gone into the development and refinement of routing architectures and algorithms for internetwork environments. In this issue we will explore some of these efforts and present routing from both a theoretical and practical perspective.

For completeness, we begin with a reprint of an overview article we ran in our June 1988 issue entitled "Routing in an Internetwork Environment" by Ross Callon, Marianne Lepp and Varda Haimo.

Jose Joaquin Garcia Luna Aceves discusses methods for avoiding routing loops in an article starting on page 8. As you will discover, these topics are not free from academic controversy; in fact Dr. Garcia Luna contradicts some of the statements made by Callon et al. As they say on Television: "Opposing views from responsible individuals are encouraged."

Much of the work to design new—or improve existing—routing algorithms is being done under the auspices of the Internet Engineering Task Force (IETF). In this issue we have articles describing two of these efforts: the OSPF Working Group and the Open Routing Working Group. The articles are by Rob Coltun and Marianne Lepp, respectively.

The *Adopt a Gateway* program was a community effort to save the Internet from terrible congestion problems in the days when all the core gateways were LSI 11s. Since then, the core gateways (or "routers" in OSI terminology) have been upgraded to BBN Butterflies. Bob Enger tells the story of how the adoption program came about and explains the infamous "Extra Hop" problem.

Routing is certainly not unique to the TCP/IP world. Paul Tsuchiya and Rob Hagens describe the emerging OSI IS-IS and ES-IS routing architectures.

It should be pointed out that this issue contains information mostly on *intra-domain* routing. Inter-domain routing is another topic which we hope to cover in future issues.

For your benefit we have included a short glossary of commonly used routing terminology, see page 25.

Finally in this issue, we bring you some announcements about current Internet activities.

The Packet Driver

by John Romkey, Epilogue Technology

What's the problem?

One of the fundamental functions of an operating system is to multiplex devices and resources among several programs that want access to them. For instance, your computer may have only one disk drive but may run many programs, perhaps simultaneously, that want to access the disk drive. Or it may run several protocols stacks (for instance, TCP/IP and Novell NetWare or Banyan Vines) that want to simultaneously access one network interface, maybe an Ethernet card.

Unfortunately, MS-DOS, the operating system most people run on their IBM PC's and clones, doesn't provide for this. In fact, it doesn't know anything about networks or network interfaces. So just about every protocol implementation has built into it a driver for the network interface it uses. Then if you run two different protocols at the same time, with their own private drivers, your PC probably crashes because the two different drivers keep interfering with one another.

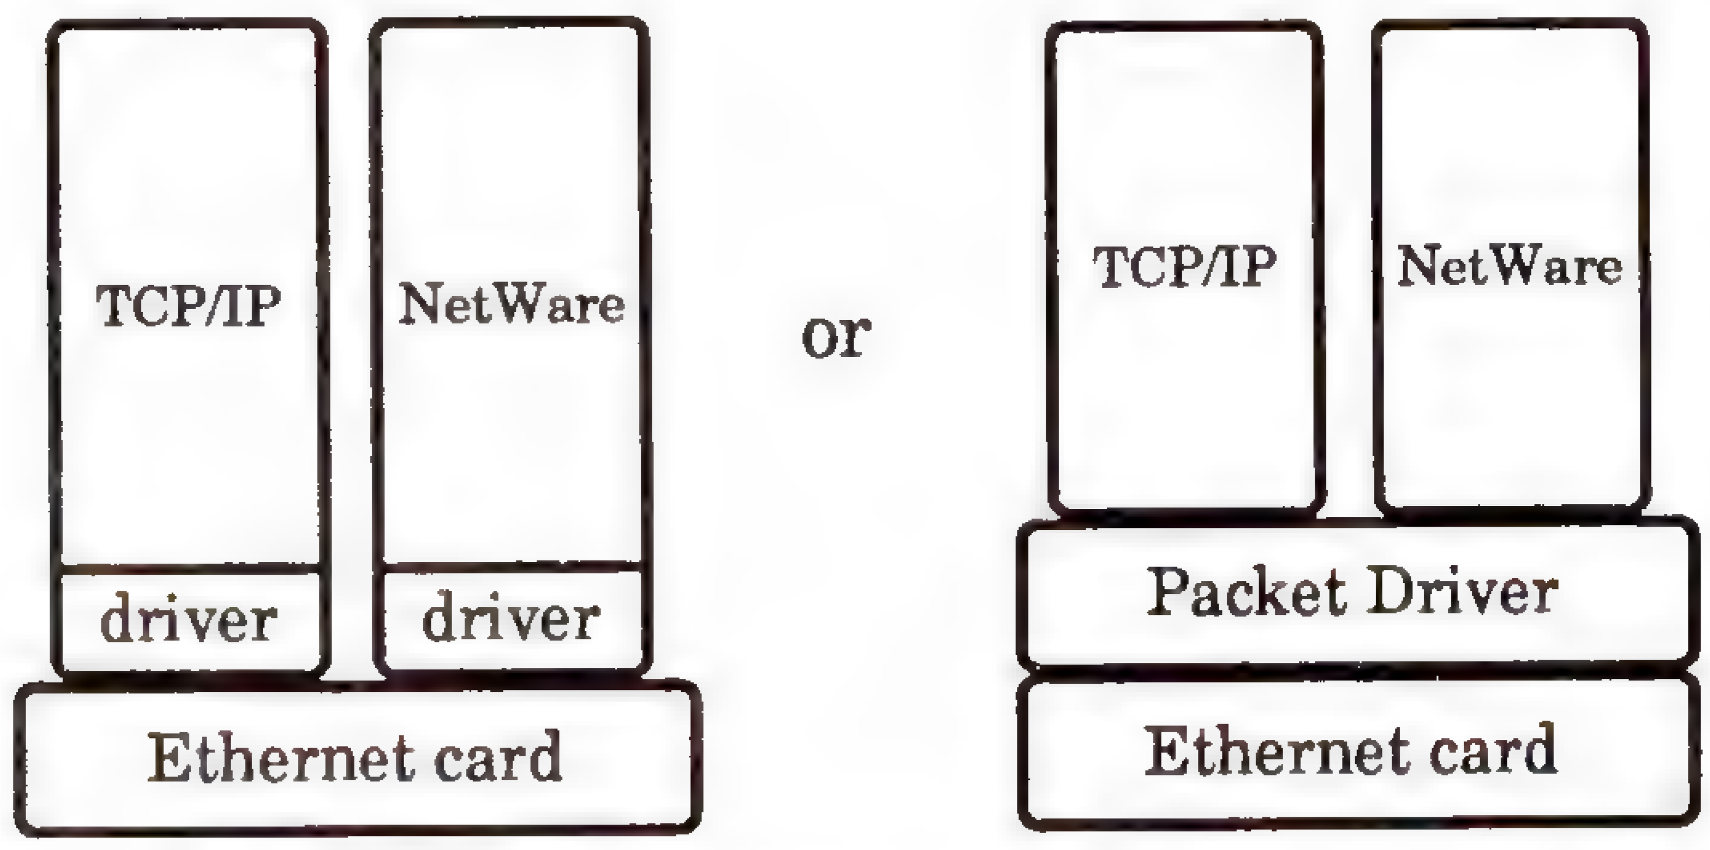

Figure 1: Conventional and Packet Drivers

Diversity of cards

Another problem is that there are at least thirty different Ethernet cards available for the IBM PC, in all shapes, sizes, colors and programming interfaces,—especially programming interfaces. The problem here is that two different cards tend to need different drivers. Unlike the world of hard disk controllers or serial cards for PC's, there is very little standardization of network cards. And Ethernet cards are only the most common; if you start counting in Token Ring cards, and ARCNET and other network interfaces, you're in real trouble.

This plethora of network cards and drivers is a software nightmare for protocol vendors, who suddenly find that they have to support every network card in the known universe, and additionally some from Andromeda, and also have to offer all these options to their customers, some of whom aren't even sure if they're using Ethernet or Token Ring, let alone which particular network card.

What's the solution?

The *Packet Driver* comes to the rescue of this situation. The Packet Driver is a specification that tells how a software driver may be loaded under DOS, hide details of the network card it operates with from protocol stacks using that card, and allow more than one protocol stack to use the card at the same time.

When the Packet Driver loads into memory, it grabs an interrupt vector (to avoid conflicts, the user may specify a particular vector).

Then it calls the MS-DOS *Terminate* and *Stay Resident* function, and the user may run more programs. Any protocol stacks which want to access the Packet Driver may then call into it via the software interrupt (they may automatically scan a range of interrupt vectors for a special signature that shows that a Packet Driver is using an interrupt, so that the user doesn't have to specify the interrupt again) and then call into it to tell it to route packets of a certain type to the stacks which want them. Since there is only one piece of software writing to the actual network card at the same time, there are no conflicts as long as different protocol stacks access different protocol types.

Each protocol stack is then linked with just a driver that knows how to access the Packet Driver loaded into the system, so there can be one version of the protocol stack rather than thirty. Optimally, the user is provided with a Packet Driver by the manufacturer of the network card they are using in their PC. Not all vendors do this yet, though, but often a Packet Driver for a particular card can be found somewhere on the Internet, for instance, in the collection of drivers maintained at Clarkson University.

What doesn't the Packet Driver do?

There are a couple of important things the Packet Driver won't do for you:

First, it won't let you run two different TCP/IPs at the same time. This is because it doesn't know anything about the protocols using it. So if you want to run two TCP/IPs at the same time, suppose a public domain one that you like because you like some special feature of the Telnet, and at the same time you also run a commercial NFS package, the Packet Driver won't help you here. The problem is that both TCP/IP's will want the same packets: IP and ARP, and the Packet Driver can only give these packets to one of them at a time.

The second problem is more subtle. A protocol stack that uses the Packet Driver doesn't need to know what particular network interface it's using, but it still needs to know the type of the network. That is, it doesn't need to know whether you're using a 3COM or a Western Digital Ethernet card, but it does need to know that the card is an Ethernet (or Token Ring or ARCNET or...) card.

IP packets are not represented the same way on all network media, and there are often differences in how IP addresses are mapped to network addresses. The same goes for other protocol stacks, like the popular proprietary PC products such as Novell *NetWare,* 3COM's *3+Open* and Banyan *Vines,* and other more open protocols such as DECnet and OSI. Since the Packet Driver is designed not to know anything about the protocols which use it, and is supposed to be simple, the task of dealing with the type of the network media falls to the protocol stack accessing the Packet Driver, and therefore it must still have dependencies on the media. Hence a generic Ethernet version, a generic Token Ring version, and others.

Some authors of Packet Drivers have gotten around this limitation by having their drivers mimic an Ethernet, by dummying up the ARP protocol and performing their own encapsulation. There exist Token Ring, SLIP and ARCNET drivers that all pretend to be Ethernet.

Similar specifications

Well after the Packet Driver specification was created, two other similar specifications showed up. The first is *NDIS*, from 3COM and Microsoft.

continued on next page

The Packet Driver *(continued)*

It has various drawbacks when compared to the Packet Driver: for instance, you cannot unload an NDIS driver or protocol stack accessing an NDIS driver, whereas you can unload it using the Packet Driver. Also, NDIS is mostly supported by 3COM cards, whereas there are Packet Drivers for many different manufacturer's network cards. Some software vendors have actually written Packet Drivers that call NDIS drivers, to allow protocol stacks that call the Packet Driver to operate over NDIS drivers.

The other specification is called *OLI*, and was introduced by Apple and Novell, obviously as a response to NDIS. Virtually no network vendors support OLI.

A little history

I originally designed the Packet Driver and wrote its specification while I was at FTP Software in 1987. After I left, James VanBokkelen maintained the specification and revised it several times with input from Packet Driver and protocol stack authors. Many freely available TCP/IP packages now support it: MIT/CMU/Harvard *PC/IP; NCSA Telnet;* and Phil Karn's *KA9Q* package. Also, most commercial packages support it: FTP Software's *PC/TCP;* Wollongong's *WIN/TCP;* and Beame & Whiteside's *BW-TCP*. Sun informally supports it with a special driver for *PC/NFS*. There exists a special (user-supported) Novell NetWare shell that calls the Packet Driver, and there are also NetWare shells that have been modified to provide a Packet Driver interface to other protocols so that you load in the NetWare shell with its built-in driver and then other protocols can access the driver by calling Packet Driver functions.

Later, Russ Nelson at Clarkson University began assembling a collection of Packet Drivers, covered by the *Free Software Foundation* "copyleft." This collection has become quite extensive, and includes Packet Drivers for many popular Ethernet cards, as well as the IBM Token Ring card, SLIP, a driver that sends IP packets over Novell's *IPX* protocol, and a driver that sends IP packets over NetBIOS. To make life simpler, these drivers masquerade as Ethernet drivers.

Resources

The Packet Driver specification is available for anonymous FTP from `vax.ftp.com`, in the file `/pub/packet-d.mss`. James VanBokkelen, who currently maintains the driver specification, may be reached at `jbvb@ftp.com`. I may be reached at `romkey@asylum.sf.ca.us`. The specification is freely usable; there are no license fees for implementing or using it, and the document is freely copyable. The Clarkson Packet Driver collection is available from host `sun.soe.clarkson.edu` by anonymous FTP. The executable files are in `/pub/packet-drivers/drivers.ar` and the source code is in `/pub/packet-drivers/driverss.ar`. Russ Nelson may be reached at `nelson@duty.clarkson.edu`

Discussions of the Packet Driver may be directed to the Internet mailing lists `tcp-ip@nic.ddn.mil` (which corresponds to `comp.protocols.tcp-ip` on USENET) or `pcip@udel.edu` (which corresponds to `comp.protocols.tcp-ip.ibmpc` on USENET) or the USENET newsgroup `comp.dcom.lans`.

JOHN ROMKEY is currently trying to live the good life in California, surrounded by friends, cats and more computers than he can count. His current projects involve creating the *Internet Toaster* and building the *Little Garden Network*.

Call for Participation

The *SIGCOMM-SIGGRAPH Workshop on Graphics and Networking* is planned for 16–18 January 1991 at the National Center for Atmospheric Research in Boulder, Colorado.

The futures of graphics and networking are closely linked together. The spread of networked workstations with high-quality bit-map displays has encouraged the creation of environments which use networks to exchange graphics images. The advent of fiber-optic networks, with gigabit data bandwidths, capable of carrying data fast enough to support real-time graphics, seems likely to encourage further integration of the two fields.

Topics

The goal of this workshop is to bring together key members of the graphics and networking communities, including people involved in research, development and standards, to discuss concerns common to the two fields. Topics of discussion are expected to include: What are the critical interactions between work in graphics and work in networking? How will current and future graphics standards affect current and future networking protocols? What effect will new developments in each field have on the other? Suggestions of additional topics are welcomed.

Format

The workshop format will be three days of plenary meetings, with each day having four sessions. Each session will start with brief presentations (5–10 minutes) by selected attendees, followed by discussion. A meeting report will appear in the *SIGGRAPH* and *SIGCOMM* newsletters.

How to participate

To ensure a good interaction between participants, the workshop is limited to no more than 75 people. We will try to give every attendee the opportunity to give a very short talk. If you are interested in attending, send a two paragraph application that describes your background, and the topic or topics you would like to discuss, to the program co-chairs by October 22, 1990:

Dr. Ralph Droms
Computer Science Dept.
Bucknell University
323 Dana Engineering
Lewisburg, PA 17837104
717-524-1145
droms@bucknell.edu

Dr. Robert Haber
Dept. of Theoretical and
Applied Mechanics
University of Illinois
111J Talbot Lab South Wright St
Urbana, IL 61801
217-333-3826
haber@ncsa.uiuc.edu

Fees

The conference fee will be $80 for SIG/ACM members and $100 for non-members. We expect a choice of hotels, with prices from $35 to $60 a night. The workshop itself will be held at the facilities of the National Center for Atmospheric Research which is generously contributing the meeting space.

Student attendance

There are spaces reserved for two graduate students to attend the workshop. These two students will be asked to keep notes, and provide a workshop report which will be printed in the *SIGGRAPH* and *SIGCOMM* newsletters. In return for this work, the workshop will provide up to $500 in travel funding for each student. Students interested in applying for these positions should contact the conference co-chairs.

A letter to the Editor

Dear Mr. Jacobsen,

I am writing to you concerning the article entitled "X Windows: More Than Just a Pretty Face," which appeared in your May 1990 edition of *ConneXions*. Although the piece is quite interesting, there is a critical error in its contents.

I was dismayed to find no mention of Hummingbird's *HCL-eXceed Family of X Servers* in the section about Software Packages.

This is especially disturbing to me since Hummingbird Communications Ltd. had a display booth at your INTEROP show in 1989, where we demonstrated our line of X Windows software.

How did you, therefore, overlook to mention us, yet at the same time include our competitors in the article?

Please let me know how you will rectify this matter.

Yours truly,

Jan Adamek,
V. P. of Marketing and Sales.
Hummingbird Communications, Ltd.

The Editor responds:

Your letter leads me to cite the fundamental, if perhaps hitherto unstated, principle of this publication:

> *ConneXions* is a *technical journal*. As such, announcements, endorsements or reviews of products (with the exception of books) are *not* presented in this forum.

To the extent that the article in question mentioned products, it was clearly for illustrative purposes, and not as a complete market survey. Indeed, if you go back and read the article again, you will find several clauses of the form "For example..." and "...to name but a few." in the context of products.

Whilst I can appreciate your dismay in not having your product mentioned when a competitor's product was, I would hardly characterize it as a "critical error." Let me re-iterate that, as a principle, we do not discuss products in this journal. The situation you complain about is purely coincidental, and not indicative of any attempt to play favorites.

It is sometimes quite tempting to share with our readers information (good or bad) about products. Every day I receive on the average 20 press releases about exciting new developments in this rapidly moving industry. However, I do feel that it is outside the current scope of *ConneXions* to engage in product analysis, and until perhaps another publication is created with such a purpose, the rules must remain as stated above.

I appreciate the feedback and hope this clarifies our editorial position.

Ole J. Jacobsen,
Editor and Publisher

Book Review

Megabit Data Communications: A Guide For Professionals by John T. Powers, Jr. and Henry H. Stair II. Published by Prentice-Hall, Inc., a Division of Simon & Schuster, 1990. ISBN 0-13-573569-6.

Audience

An interesting book explaining T-Carrier in some detail has recently been published entitled, "Megabit Data Communications." Subtitled, "A Guide For Professionals," this book describes practical applications of megabit-speed digital transmission technologies, products, and services. It is directed at managers, engineers, planners and designers who deal directly with digital communications.

Everything in one place

The authors, Jack Powers and Pete Stair, note that the book resulted from their reflections on the data communications business and seeing the surprising difficulty which even simple tasks require. When it became apparent to them that information needed to plan, specify, engineer and install high-speed facilities was spread thinly over a variety of sources—some of which were quite obscure—they decided to write this book and bring the information together in one place.

Content

They do not discuss prices, delivery or vendor performance, simply because such information would be obsolete before the book was published. What they do discuss in detail includes:

- ISDN networks
- T-Carrier services and related hardware
- AT&T's Dataphone digital services
- Telex and TWX
- Voice technologies
- Fiber optic transmission techniques
- Private digital services
- Multivendor integration

There are numerous charts, diagrams, drawings and other illustrations to assist in understanding what they have written.

Valuable and useful

You might find this book to be a valuable and useful addition to your telecom library. It certainly will assist in evaluating vendor's claims as to equipment performance and compatibility.

It should be available at this time in the technical department of bookstores in your community. —*Patrick Townson*

CONNEXIONS
480 San Antonio Road
Suite 100
Mountain View, CA 94040
415-941-3399
FAX: 415-949-1779

Bulk Rate
U.S. POSTAGE
PAID
SAN JOSE, CA
PERMIT NO. 1

FORWARDING AND ADDRESS CORRECTION REQUESTED